arXiv:1408.0840v1 [gr-qc] 5 Aug 2014

Detecting compact binary coalescences with seedless clustering

M Coughlin,[1, a] E Thrane,[2] and N Christensen[3]

[1]*Department of Physics, Harvard University, Cambridge, MA 02138, USA*
[2]*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, CA, 91125, USA*
[3]*Physics and Astronomy, Carleton College, Northfield, Minnesota 55057*

(Dated: August 6, 2014)

Compact binary coalescences are a promising source of gravitational waves for second-generation interferometric gravitational-wave detectors. Although matched filtering is the optimal search method for well-modeled systems, alternative detection strategies can be used to guard against theoretical errors (e.g., involving new physics and/or assumptions about spin/eccentricity) while providing a measure of redundancy. In previous work, we showed how "seedless clustering" can be used to detect long-lived gravitational-wave transients in both targeted and all-sky searches. In this paper, we apply seedless clustering to the problem of low-mass ($M_{\rm total} \leq 10M_\odot$) compact binary coalescences for both spinning and eccentric systems. We show that seedless clustering provides a robust and computationally efficient method for detecting low-mass compact binaries.

PACS numbers: 95.75.-z,04.30.-w

I. INTRODUCTION

Compact binary coalescences (CBCs) of black holes (BHs) and/or neutron stars (NSs) are a likely source of gravitational waves (GWs) [1–3]. CBC events include binary neutron stars (BNSs), neutron-star black holes (NSBHs), and binary black holes (BBHs). As a CBC passes through its inspiral and merger stage, it generates GWs which sweep upward in frequency and strain amplitude through the sensitive band of GW detectors. The detection of GWs from CBCs will provide information about the populations of compact objects in the universe, elucidate the properties of strong field gravity, and provide a means to test general relativity.

Here, we focus on relatively low-mass binaries ($M_{\rm total} \leq 10M_\odot$). There are two reasons for restricting our attention to this region of parameter space. First, the rate of low-mass CBCs is less subject to theoretical uncertainty than high stellar-mass binary BHs and intermediate-mass BBH. Second, we are interested in long-lived signals ($\approx$ 54–270 s), which appear as curved tracks in spectrograms of GW strain power, and therefore provide an appealing target for seedless clustering [4, 5] (described in greater detail below).

Searches for CBCs often use matched filtering, which requires precise knowledge of astrophysical waveforms. (Excess power searches are also used, especially for high-mass systems associated with shorter signals; see, e.g., [6].) Since CBCs are, for the most part, well-modeled systems, matched filtering provides an essentially optimal strategy for detecting compact binaries. However, there are several reasons why it is useful to consider alternative detection strategies.

Independent verification. Alternative methods can provide independent verification of detections by matched filter pipelines, thereby increasing confidence in the veracity of a result (of course, because seedless clustering will be less sensitive than matched filtering searches in most cases, non-detection by seedless clustering is not a concern either). Although there is some redundancy provided by the multiple implementations of matched filtering used in current searches, seedless clustering provides a very different approach to gravitational-wave detection and detection by both methods potentially indicates the robustness of the result.

Visualization of the GW signal. In general, advanced detector CBC events are expected to be buried in noise to the extent that it will be difficult to see by eye their signature in a time series or strain auto-power spectrogram. Here we show that, by coherently combining the output of multiple detectors, CBCs can be visualized as faint but visible arcs on "radiometric spectrograms"—especially when the eye is guided by the reconstructed track of a search algorithm. (The curious reader is encouraged to skip ahead to Fig. 1 for an example of a radiometric spectrogram.)

Visualizing the signal helps confirm that the detected signal looks like one expects. The coherent combination also allows for confirmation that the parameter estimation of the signals, including the direction, masses, and time of coalescence, are all consistent "by eye" with the radiometric spectrogram. For example, an error in the reconstructed CBC direction creates characteristic stripes; see [5]. If, on the other hand, the masses are incorrect, the reconstructed track will have the wrong frequency evolution as a function of time.

Data processing corner cases. Real-world GW searches require design choices, which take into account the complicated nature of GW detectors. Detector performance is non-stationary, the noise contains non-Gaussian "glitches," and data-taking is sometimes interrupted by lock-loss, just to name a few relevant effects. As a result, workarounds are employed, e.g., to estimate background, to discard noisy data, and to handle gaps. Matched fil-

[a]Electronic address: coughlin@physics.harvard.edu

tering [7] and seedless clustering [8] have different ways of performing these tasks. In general, these technical details are (by design) not important factors in determining the average sensitivity of a search. However, by employing multiple search methods we can guard against individual events falling between the cracks.

An example of a possible data processing corner case is shown in Fig. 2. This event was identified correctly by both matched filtering and seedless clustering. The left-hand panel shows $\rho(t;f)$ obtained using "engineering-run" data[34], in which data from a LIGO sub-system, in this case the pre-stabilized laser, is recolored to match the Advanced LIGO noise curve. Such engineering run data does not contain astrophysically useful strain measurements, but is nonetheless useful for its non-Gaussian noise characteristics. The data contains five segments consistent with non-Gaussian noise and would be removed in a search. Nonetheless, despite these noise artifacts, it is still possible to detect a simulated binary neutron star signal. The right-hand panel shows the reconstructed signal, obtained with the seedless clustering algorithm we describe below.

Waveform uncertainties. Theoretical errors in matched filter waveforms can arise from the computational limitations and/or imperfect approximations. Due to computational limitations, most CBC searches so far use template banks composed of non-spinning, non-eccentric waveforms, which are less computationally challenging than search with spin and eccentricity. High-spin systems take longer to simulate with numerical relativity and the addition of extra spin parameters creates larger, more unwieldy template banks. Searches that ignore spin can suffer significant losses in sensitivity [9]; the cases considered here have between a 23-36% match, which is maximized over time and phase, between the spinning and non-spinning waveforms; these numbers increase to at least 97% when maximized over mass as well. When spin is included, it is often assumed that the spins are aligned in order to make the calculation more tractable. Even so, the inclusion spin effects can lead to a factor of two increase in sensitive volume [10]; see also [11, 12]. Main sequence binaries circularize by the time they enter the sensitive frequency band of terrestrial detectors [13]. However, dynamical capture may produce gravitational waves from highly eccentric binaries [13–17]. The cases considered here have a less than 1% match between the eccentric and circular waveforms of equivalent mass. These numbers increase to between 20-60% when maximized over mass. This highlights the difficulty of detecting them using a template bank composed of circular templates. Imperfect assumptions about eccentricity and spin may therefore create openings for ostensibly sub-optimal detection strategies.

New physics. One can also imagine significant waveform errors due to the existence of new or unforeseen physics. For example, Piro raised concerns about the effects of magnetic interactions in BNS [18]. While these magnetic interactions were subsequently shown to be ignorable [19] for GW astronomy, one can imagine a comparable source of theoretical error. More speculatively, non-standard theories of gravitation can lead to modifications of the waveform [20, 21].

Thus, there are many reasons why it is worth considering alternatives to matched filtering. One common alternative technique for detection of GW transients is to search for excess power in spectrograms (also called frequency-time ft-maps) of GW detector data [22–24]. This method casts GW searches as pattern recognition problems.

Previous work has shown how "seedless clustering" can be used to perform sensitive searches for long-lived transients [4, 5]. The idea of seedless clustering is to integrate the signal power along spectrogram tracks chosen to capture the salient features of a wide class of signal models. Seedless clustering calculations are embarrassingly parallel, and so the technique benefits from the recent proliferation of highly parallel computing processors including graphical processor units and multi-core central processing units. Previous papers [4, 5] have pointed out that seedless clustering algorithms might be useful for CBC detection/confirmation.

In this work, we apply the seedless clustering formalism to efficiently search for CBC signals. In section II, we review the basics of seedless clustering. We show how the formalism of [4, 5] can be tuned to more sensitively detect CBC signals. In section IV, we determine the sensitivity of seedless clustering algorithms (with different levels of tuning) to CBC waveforms. We conclude with a discussion of topics for further study in section V.

II. SEEDLESS CLUSTERING FOR CHIRPS

Searches for unmodeled GW transients typically begin with spectrograms proportional to GW strain power. The pixels of these spectrograms are computed by dividing detector strain time series in segments and computing Fourier transform of the segments. The Fourier transform of the strain data from detector I for the segment with a mid-time of t is denoted $\tilde{s}_I(t;f)$. For the results presented here, we use 50%-overlapping, Hann-windowed segments with duration of 1 s. The frequency resolution is 1 Hz.

Searches for long-duration GW transients in particular use the cross-correlation of two GW strain channels from spatially separated detectors to construct ft-maps of cross-power signal-to-noise ratio [24]:

$$\rho(t;f|\hat{\Omega}) = \mathrm{Re}\left[\lambda(t;f)e^{2\pi i f\Delta\vec{x}\cdot\hat{\Omega}/c}\tilde{s}_I^*(t;f)\tilde{s}_J(t;f)\right]. \quad (1)$$

Here, $\hat{\Omega}$ is the direction of the GW source, $\Delta\vec{x}$ is a vector describing the relative displacement of the two detectors, c is the speed of light, and $e^{2\pi i f\Delta\vec{x}\cdot\hat{\Omega}/c}$ is a direction-dependent phase factor, which takes into account the time delay between the two detectors. The $\lambda(t;f)$ term is

a normalization factor, which employs data from neighboring segments to estimate the background at time t:

$$\lambda(t;f) = \frac{1}{\mathcal{N}}\sqrt{\frac{2}{P_I'(t;f)P_J'(t;f)}}. \quad (2)$$

$P_I'(t;f)$ and $P_J'(t;f)$ are the auto-power spectral densities for detectors I and J in the segments neighboring t. For additional details, see [4, 5, 24]

GWs appear as tracks or blobs in the ft-maps. The morphology of the the GWs are dependent on the signal. CBC signals appear as chirps of increasing frequency. Clustering algorithms are used to identify clusters of pixels Γ likely to be associated with a GW signal. The total signal-to-noise ratio for a cluster of pixels can be written as a sum of over $\rho(t;f|\hat{\Omega})$:

$$\mathrm{SNR}_{\mathrm{tot}} \equiv \frac{1}{N^{1/2}} \sum_{\{t;f\}\in\Gamma} \rho(t;f|\hat{\Omega}), \quad (3)$$

where N is the number of pixels in Γ.

Different clustering algorithms employ different methods for choosing Γ. Seed-based algorithms connect statistically significant seed pixels to form clusters [4, 25]. In seedless clustering algorithms [4], Γ is chosen from a bank of parametrized frequency-time tracks. Each such track is referred to as a "template." Calculations for many templates can be carried in parallel, which facilitates rapid calculations on multi-core devices such as graphical processor units (GPUs).

It must be noted that our templates are different from matched-filtering templates. Seedless clustering templates describe the morphology of a power spectrogram track whereas matched filter templates describe the phase evolution of a signal appearing in just one detector. Matched filter templates contain all the available information about the signal, whereas seedless clustering templates are a lossy, binned representation of the signal. By throwing away information, the seedless clustering search is less sensitive than a matched filter search, but by throwing away information, it can simultaneously become more robust against waveform uncertainties and new physics.

A very general search with minimal assumptions may employ, e.g., a template bank of randomly generated Bézier curves [26], which have been shown to do a reasonably good job of mimicking long-lived narrowband gravitational-wave signals [5]. However, one may equally well carry out a more specialized search, targeting a specific class of signals. Given our present interest in CBC signals, we opt to work with a more specialized template bank consisting of parametrized chirps:

$$f(t) = \frac{1}{2\pi}\frac{c^3}{4GM_{\mathrm{total}}}\sum_{k=0}^{7} p_k \tau^{-(3+k)/8}, \quad (4)$$

where

$$\tau = \frac{\eta c^3 (t_c - t)}{5GM}. \quad (5)$$

Here, G is the gravitational constant and M_{total} is the total mass of the binary. The expansion coefficients p_k can be found in [27]. Each chirp template is parametrized by two numbers: the coalescence time and the chirp mass. (This is in contrast to Bézier curves, which are parametrized by six numbers.) While technically, the waveform depends on the individual component masses, the main features of the signal can be well-approximated by only the chirp mass. Therefore, to reduce the parameter space by one variable, we employ the approximation that the individual component masses are equal.

The space of arbitrary long-lived gravitational-wave signals is very large, and so general algorithms, employing Bézier templates, use randomly generated numbers to span as much of the signal space as possible. The space of CBC chirp signals is much smaller. A search for binary neutron star signals with component masses of 1.4–3$M_\odot$ plateaus in sensitivity using just 50 chirp mass bins. The same search requires 825 time bins for 660 s of data—a typical on-source window (in which the signal is assumed to exist) for a GW search triggered by a gamma-ray burst [28]. Thus, the template density for a targeted CBC search is $\approx 6.3\times10^4\,\mathrm{ks}^{-1}$ whereas a Bézier bank might require $\approx 8\times10^8\,\mathrm{ks}^{-1}$ [5]; (1 ks = 1000 s). Since there are so many fewer templates in a chirp-template search (about four orders of magnitude fewer), it is computationally feasible to employ every template.

Of course, the previous calculation was for a targeted search in which the source location is previously determined, e.g., by an electromagnetic trigger or by a different search algorithm. In [5], we showed how the introduction of a phase factor can be used to allow to carry out an efficient all-sky search with seedless clustering. This formalism is straightforwardly applied to our chirping templates. For the CBCs detected by the LIGO Hanford-Livingston detector pair, it is sufficient to consider 40 time delays, each corresponding to a ring on the sky. Thus, even an all-sky search using seedless clustering to detect CBC signals can employ a relatively modest template density: $\approx 2.5\times10^6\,\mathrm{ks}^{-1}$.

To estimate the computational cost of an all-sky seedless clustering search (with chirp-like templates), we carried out a benchmark study using a Kepler GK104s GPU and an 8-core Intel Xeon E5-4650 CPU. Each job was allotted 8 g of memory. The GPU was able to analyze 660 s of data in 48 s, corresponding to a duty cycle of $\approx 7\%$. Using all eight cores, the CPU duty cycle was comparable; the job-by-job variability in run time is greater than the difference between GPUs and 8-core CPUs on average.

If we require background estimation at the level of FAP = 1%, it follows that a continuously running seedless clustering search with chirp-like templates can be carried out with just 8 GPUs (or 8-core CPUs). (Background estimation at the level of FAP = 0.1% would require 74 GPUs / 8-core CPUs.) In reality, the duty cycle from coincident GW detectors may be $\approx 50\%$, in which case these computing requirements are conservative by a

factor of two. Repeating the test for a targeted search (for which the source location is known), we obtained an 8-core CPU duty cycle of 2%, a factor of 3 speed-up. The targeted search run on GPUs does not run appreciably faster than the all-sky version.

In addition to improved computational efficiency, there is another important advantage to be gained through the use of CBC templates compared to Bézier templates. In [5], we showed that quadratic Bézier curves do a mediocre job approximating CBC signals. By adopting the chirping templates described in Eq. 4, we expect to capture more signal-to-noise ratio, and thereby extend the sensitive distance of the search. An example of a weak BNS signal recovered with a seedless chirping template is shown in Fig. 1.

III. SENSITIVITY STUDY

In order to determine the sensitivity of seedless clustering with chirp templates, we perform a sensitivity study with Monte Carlo noise. We assume Gaussian noise consistent with the design sensitivity of Advanced LIGO. Following [28], we assume that an external trigger has predicted that the signal exists in a 660 s on-source window. For each trial, we search for a chirp signal four different ways: using Bézier curves and a known sky location (BK), using Bézier with unknown sky location (BU), using chirp templates with a known sky location (CK), and using chirp templates with an unknown sky location (CU).

The first step of the sensitivity study is background estimation. We perform many trials to estimate the distribution of $\mathrm{SNR}_{\mathrm{tot}}$ for noise. We generate separate noise distributions for all four search variations (BK, BU, CK, CU). Using these noise distributions, we determine the value of $\mathrm{SNR}_{\mathrm{tot}}$ (for each search variation), which corresponds to a false alarm probability (FAP) of 0.1%.

The next step is to determine the distance to which different signals can be detected with $\mathrm{SNR}_{\mathrm{tot}}$ sufficient for a detection with FAP $< 0.1\%$. We add GW signals to realizations of detector noise. Each injected signal is injected with an optimal sky location and an optimal source orientation. We define the "sensitive distance" as the distance at which 50% of the singals are recovered with FAP $< 0.1\%$. We consider 14 CBC waveforms with component masses ranging from 1.4–3$M_{\odot}$. Of these waveforms, eight characterize eccentric systems and three characterize systems where one or more object has a large dimensionless spin:

$$a \equiv cJ/Gm^2. \tag{6}$$

Here J is the angular momentum and m is the component mass.

Non-eccentric waveforms are generated using a SpinTaylorT4 approximation. Eccentric waveforms are generated using `CBWaves`, which employs all the contributions that have been worked out for generic eccentric orbits up to 2PN order [29]. The parameters for each waveform are give in Table I.

IV. RESULTS

The results of our sensitivity study are summarized in Table I. There are a number of interesting trends. First, while all of the (known direction) seedless clustering distances are astrophysically interesting, the chirping templates perform consistently better than the Bézier templates. The ratio of detection distance for chirping templates / Bézier templates ranges from 100–161% with a mean of 131%. The average ratio of sensitive volumes is 240%. This is comparable to the gain in sensitivity for a matched filter search to be had through the inclusion of spin [10], indicating that while it is certainly advantageous to use chirping templates, the Bézier templates do surprisingly well.

We find no significant difference in the chirping-template detection distance between systems that do or do not contain spin. The similarity in the sensitivity distances between the non-spinning and spinning cases indicates that the spins do not effect the signal morphology in a significant enough way to deviate from the non-spinning track. The advantage of chirping templates appears to increase slightly for spinning systems. Finally, we observe no loss in sensitivity going from the targeted CK search to the all-sky CU search. Evidently, the increase in signal space from the additional parameter of sky location is not sufficient to meaningfully affect the background distribution.

Highly eccentric signals generally have a longer duration than than those with low or no eccentricity. The sensitivity (using both Bézier templates and chirp-like templates) decreases with eccentricity. There is a similar advantage in distance of the chirp-like templates over the Bézier templates at low eccentricity. This benefit decreases slightly as eccentricity increases. This is due to the breakdown of the circular binary approximation. The breakdown becomes more pronounced at higher eccentricities, as one would expect.

Using matched filtering, Advanced LIGO, operating at design sensitivity, is expected to reliably detect BNS (with optimal orientation and sky location) out to distances of 450 Mpc [30], 2.4× further than the seedless clustering detection distance quoted here. It follows that $\approx 8\%$ of the events detected by matched filtering will produce a FAP $< 0.1\%$ signature when followed up with seedless clustering. Given a realistic astrophysical rate of $40\,\mathrm{yr}^{-1}$ BNS detections by Advanced LIGO [31], this implies that we can expect to confirm ≈ 3 events per year of science data using seedless clustering. The use of expanded template banks that include waveforms with spin will allow matched filtering searches to observe to comparable distances as the non-spinning case [11]; therefore the follow-up numbers will be similar to the above. In the event that circular template banks are used to search

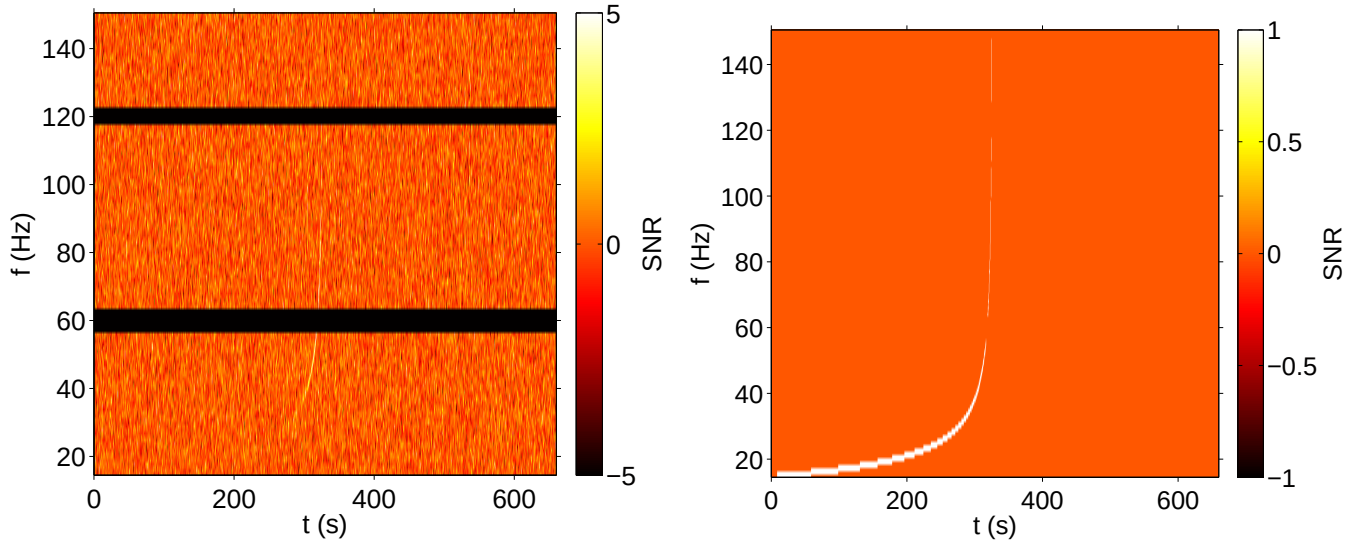

FIG. 1: The plot on the left shows $\rho(t;f)$ for a simulated eccentric ($\epsilon = 0.2$) BNS signal injected on top of Monte Carlo detector noise. The component masses are $1.4M_{\odot}$. The chirping signal appears as a faintly-visible track of lighter-than-average pixels. The horizontal lines are frequency notches to remove instrumental artifacts. On the right is the recovery obtained with seedless clustering. The signal is recovered with a FAP $< 0.1\%$.

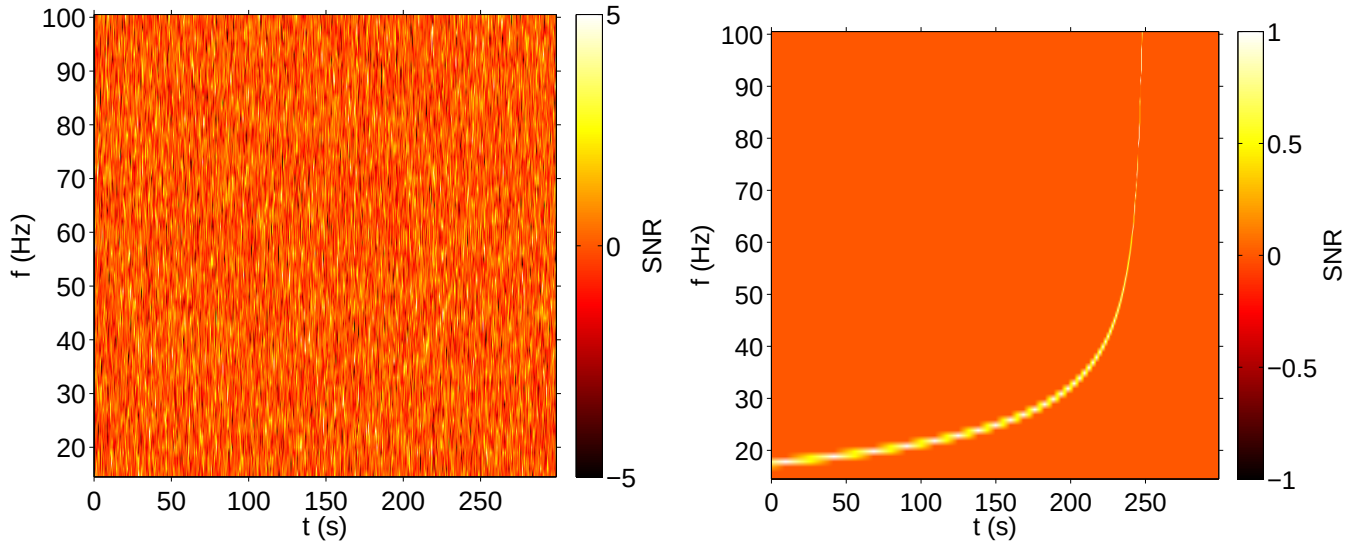

FIG. 2: The plot on the left shows $\rho(t;f)$ during a recent LIGO engineering run, in which data from a LIGO subsystem—not sensitive to GW strain—is recolored to produce semi-realistic detector noise. On the right is the seedless clustering recovery, which is able to to detect the injected signal with high confidence FAP $< 0.1\%$, despite the relatively poor data quality. (Though it is not immediately apparent from these plots, five segments are identified as characteristic of non-stationary noise [8].) The signal is recovered with FAP $< 0.1\%$.

for eccentric signals, there will be a non-negligable loss in sensitivity for these searches. Huerta and Brown estimate signal-to-noise ratio loss factors of about 0.5 and 0.2 for BNS systems with eccentricities of 0.2 and 0.4 respectively [32]. This would bring the matched filtering sensitivity distances of these signals to 225 Mpc and 90 Mpc; therefore seedless clustering may provide further opportunities for observing these types of signals.

In addition to providing confirmation of these loudest CBC events, seedless clustering will provide a safety net by potentially detecting events missed due to waveform error, data-processing subtleties, and/or new physics.

V. CONCLUSIONS

Seedless clustering provides a computationally efficient tool for the follow-up and detection of compact binary

Waveform	m_1	m_2	a_1	a_2	ϵ	$t_{\rm dur}$ (s)	$D_{\rm BK}$	$D_{\rm BU}$	$D_{\rm CK}$	$D_{\rm CU}$
BNS 1	1.4	1.4	0	0	0	170	160	130	190	190
NSBH 1	3.0	1.4	0	0	0	96	200	200	290	290
NSBH 2	3.0	1.4	0.95	0	0	97	220	180	320	320
BBH 1	3.0	3.0	0	0	0	54	330	330	470	470
BBH 2	3.0	3.0	0.95	0	0	55	320	270	470	470
BBH 3	3.0	3.0	0.95	0.95	0	55	320	260	470	470
BBH 4	5.0	5.0	0	0	0	42	620	560	750	750
BBH 5	5.0	5.0	0.95	0	0	42	680	680	750	750
BBH 6	5.0	5.0	0.95	0.95	0	43	750	680	830	830
EBNS 1	1.4	1.4	0	0	0.2	120	150	120	180	180
EBNS 2	1.4	1.4	0	0	0.4	224	150	120	160	160
ENSBH 1	3.0	1.4	0	0	0.2	69	180	180	290	290
ENSBH 2	3.0	1.4	0	0	0.4	127	180	160	240	240
ENSBH 3	3.0	1.4	0	0	0.6	237	180	160	240	240
EBBH 1	3.0	3.0	0	0	0.2	40	270	220	320	320
EBBH 2	3.0	3.0	0	0	0.4	70	220	200	240	240
EBBH 3	3.0	3.0	0	0	0.6	128	220	180	220	220

TABLE I: Sensitive distances for different waveforms (assuming optimal sky location and source orientation) given the design sensitivity of Advanced LIGO [30]. Each row represents a different waveform: BNS="binary neutron star," NSBH="neutron star black hole binary," BBH="binary black hole. A waveform beginning with an "E" is eccentric. The columns marked m_1 and m_2 give the component masses in units of $M_\odot$. The columns marked a_1 and a_2 give the component spins; see Eq. 6. The next columns list the ellipticity ϵ and the waveform duration in seconds. The final four columns list the (FAP = 0.1%, FDP = 50%) detection distance (in Mpc) for Bézier templates with known sky location (BK), Bézier templates with unknown sky location (BU), chirp-like templates with known sky location (CK), and chirp-like templates with unknown sky location (CU).

coalescences. While seedless clustering is expected to be less sensitive than matched filtering, it provides a number of useful features including independent verification, visualization of the gravitational-wave signal, the ability to catch corner-case signals, and robustness to both waveform uncertainty and existence of new physics.

We compared a specially tuned implementation of seedless clustering, optimized for compact binary coalescences, to a more generic search using Bézier curves. We find that the CBC-tuned search can expand the sensitive volume by as much as a factor of 4.2× depending on the waveform (a factor of two on average) compared to the generic Bézier search. Perhaps more importantly, the tuned search requires 10^4 fewer templates per unit of time, allowing for a significantly faster search.

There are a number of potential improvements to the algorithm worth exploring. It may be possible to improve the implementation of seedless clustering described here by more optimally weighting different time-frequency bins based on the known waveform. It is also worth exploring the effect of only using equal mass templates to recover potentially non-equal mass signals. It is possible that for cases with larger mass ratios than those considered here, it may be necessary to relax this assumption in order to reconstruct a majority of the signal. Another possibility for improvement is the implementation of a better parametrization for the eccentric waveforms. Unlike for circular binaries, there is, at present, no closed expression for the phase evolution of a binary with arbitrary eccentricity. A more effective parametrization is likely to capture more signal-to-noise ratio, thereby extending the sensitive range. Finally, it will be useful to carry out a systematic comparison of seedless clustering with matched filtering pipelines and using non-Gaussian noise.

VI. ACKNOWLEDGMENTS

MC is supported by National Science Foundation Graduate Research Fellowship Program, under NSF grant number DGE 1144152. ET is a member of the LIGO Laboratory, supported by funding from United States National Science Foundation. NC's work was supported by NSF grant PHY-1204371. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This paper has been assigned LIGO document number P1400110-v5.

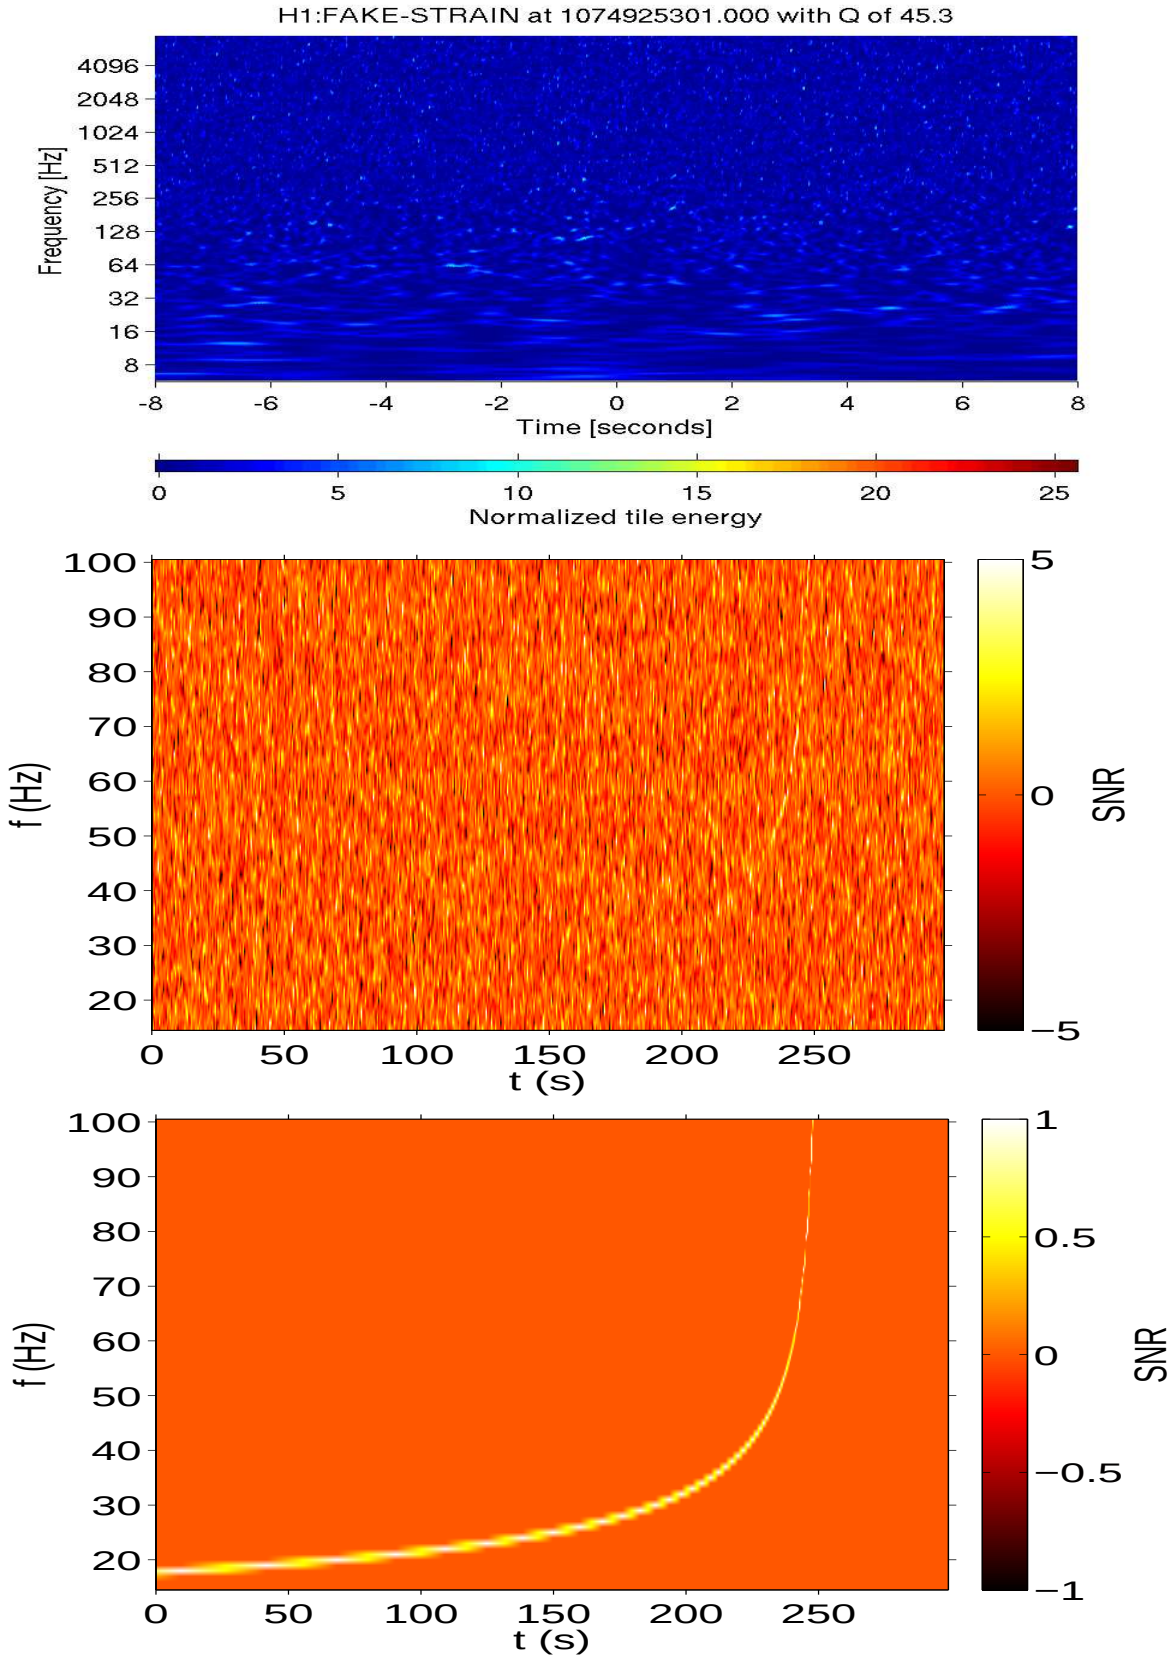

FIG. 3: Data from a recent LIGO engineering run, in which data from a LIGO subsystem—not sensitive to GW strain—is recolored to produce semi-realistic detector noise. A simulated binary neutron star signal has been added to the data. The top plot shows a wavelet transform of single-detector auto-power [33], currently in wide use for diagnostics. The injected waveform, which ends at $t = 0$, is difficult to make out by eye. The middle plot shows a spectrogram of $\rho(t; f)$. The injection, though faint, is visible between 200-250 s. The bottom plot shows the reconstructed track using seedless clustering; FAP $< 0.1\%$.

[1] J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **87**, 022002 (2013).

[2] J. Abadie et al. (LIGO Scientific Collaboration and Virgo Collaboration), Physical Review D **85**, 082002 (2012).

[3] Abadie J. et al. (LIGO Scientific Collaboration and Virgo Collaboration), Classical and Quantum Gravity **27**, 173001 (2010).

[4] E. Thrane and M. Coughlin, Phys. Rev. D **88**, 083010 (2013).

[5] E. Thrane and M. Coughlin, Phys. Rev. D **89**, 063012 (2014).

[6] J. Abadie et al., Phys. Rev. D **85**, 102004 (2012).

[7] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, K. Cannon, C. D. Capano, J. H. Clayton, T. Cokelaer, J. D. E. Creighton, T. Dent, et al., Phys. Rev. D **87**, 024033 (2013), URL `http://link.aps.org/doi/10.1103/PhysRevD.87.024033`.

[8] T. Prestegard, E. Thrane, et al., Class. Quant. Grav. **29**, 095018 (2012).

[9] Canton T. et al, *Implementing a search for aligned-spin neutron star – black hole systems with advanced ground based gravitational wave detectors* (2014), URL `http://arxiv.org/abs/1405.6731`.

[10] S. Privitera, *The importance of spin for observing gravitational waves from coalescing compact binaries with LIGO and Virgo.* (2014).

[11] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, Phys. Rev. D **86**, 084017 (2012).

[12] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, Phys. Rev. D **89**, 024003 (2014).

[13] D. A. Brown and P. J. Zimmerman, Phys. Rev. D **81**, 024007 (2010).

[14] R. M. O'Leary, B. Kocsis, and A. Loeb, MNRAS **395**, 2127 (2009).

[15] U. Sperhake, E. Berti, V. Cardoso, J. A. González, B. Brügmann, and M. Ansorg, Phys. Rev. D **78**, 064069 (2008).

[16] I. Hinder, F. Herrmann, P. Laguna, and D. Shoemaker, Physical Review D. **82** (2010), ISSN 0556-2821.

[17] R. Gold, S. Bernuzzi, M. Thierfelder, B. Brugmann, and F. Pretorius, Phys.Rev. **D86**, 121501 (2012), 1109.5128.

[18] T. Piro, Astrophys. J. **755**, 80 (2012).

[19] D. Lai, **757**, L3 (2012).

[20] M. Shibata, K. Taniguchi, H. Okawa, and A. Buonanno, Phys. Rev. D **89**, 084005 (2014), URL `http://link.aps.org/doi/10.1103/PhysRevD.89.084005`.

[21] E. Barausse, C. Palenzuela, M. Ponce, and L. Lehner, Phys. Rev. D **87**, 081506 (2013), URL `http://link.aps.org/doi/10.1103/PhysRevD.87.081506`.

[22] P. Sutton et al., New Journal of Physics **12**, 053034 (2010).

[23] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, Class. Quant. Grav. **25**, 114029 (2008).

[24] E. Thrane et al., Physical Review D **83**, 083004 (2011).

[25] T. Prestegard and E. Thrane, *Burstegard: a hierarchical clustering algorithm* (2012), `https://dcc.ligo.org/LIGO-L1200204-v1`.

[26] G. Farin, *Curves and Surfaces for CAGD, Fourth Edition: A Practical Guide* (Academic Press, 1996).

[27] Blanchet L., Living Reviews in Relativity **4** (2006), URL `http://www.livingreviews.org/lrr-2006-4`.

[28] Abadie J. et al. (LIGO Scientific Collaboration and Virgo Collaboration), The Astrophysical Journal **760**, 12 (2012).

[29] Csizmadia P. et al., Classical and Quantum Gravity **29**, 245002 (2012).

[30] J. Aasi et al., *Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories* (2013), URL `http://arxiv.org/abs/1304.0670`.

[31] J. Abadie et al., Class. Quant. Grav **27**, 173001 (2010).

[32] E. A. Huerta and D. A. Brown, Phys. Rev. D **87**, 127501 (2013).

[33] J. Rollins, *Multimessenger Astronomy with Low-Latency Searches for Transient Gravitational Waves.* (2011).

[34] The data in question are from LIGO Engineering Run 5, collected during 2014.